Office de la propriété intellectuelle du Canada

Un organisme d'Industrie Canada

Canadian Intellectual Property Office

An Agency of Industry Canada

Bureau canadien des brevets

Certification

Canadian Patent Office

Certification

La présente atteste que les documents ci-joints, dont la liste figure ci-dessous, sont des copies authentiques des documents déposés au Bureau des brevets.

This is to certify that the documents attached hereto and identified below are true copies of the documents on file in the Patent Office.

Specification and Drawings, as originally filed, with Application for Patent Serial No: **CA 2314665**, on July 28, 2000, by **CATENA NETWORKS CANADA INC.**, assignee of Francois Tremblay and Scott McClennon, for "Loop and Rate-Dependent DS Power Cutback for Power Savings in a XTU-C DSL Transceiver".

L. Réginbal

Agent certificateur/Certifying Officer

March 31, 2009

Date

Canada

(CIPO 68)
07-10-08

OPIC CIPO

LOOP- AND RATE-DEPENDENT DS POWER CUTBACK FOR POWER SAVINGS IN A XTU-C DSL TRANSCEIVER

Background of Invention:

The challenge solved with this invention is the reduction of xDSL transceiver power requirements on subscriber loops with excess SNR margin for a given required link capacity.

ADSL ATU-C transmitters reduce their transmit power on very short loops to avoid overloading the remote ATU-R receiver. This "politeness cutback" is determined according to an estimate of loop loss in the upstream signal path and applies on loops shorter than ~2-3kft 26AWG-equivalent. Other DS transmitter power cutbacks are also specified in the splitterless ADSL draft standard, for the purpose of reducing the DS signal level at the remote end in the presence of an off-hook telephone set. There is no agreed mechanism to implement DS power cutback for other purposes; specifically, there is no mechanism defined for power cutback, to reduce the ATU-C power requirements, where there is excess SNR margin at the ATU-R receiver.

The benefit of such a cutback capability can be significant in Digital Loop Carrier applications, for example, where subscriber loops are typically shorter than seen at CO-resident line interfaces and where power and thermal budgets are also very tight. Assuming a DMT transmit signal with peak-to-average-ratio (PAR) of 14.5dB[1], DS power cutback can yield significant savings on line driver AC power alone, with fixed driver supply rails. As illustrated below, these savings can exceed 500mW for a class ? driver with 13dB of cutback.

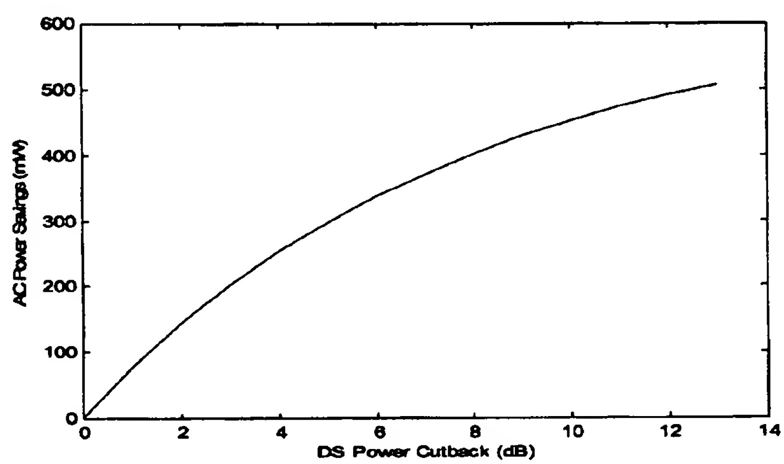

Figure 1: Fixed Supply Line Driver Power Savings vs. DS Power Cutback

[1] PAR reduction techniques claiming 3-5dB reductions exist in the literature. In practice, however, these techniques lose their effectiveness when the signal passes through analog or digital filters in the transmit path.

Figures 2-4 illustrate the attainable DS rates[2] over loop lengths of 0-12kft (26AWG), with DS power cutback to 15dB in the presence of crosstalk from other DSLs on adjacent pairs. The crosstalk scenario for the three figures is 49 other G.lite (splitterless ADSL) disturbers, 10 HDSL disturbers or 5 adjacent binder T1 disturbers, respectively.

The DS Politeness Cutback (of 0-12dB) is not shown here (will apply at loop lengths up to ~3kft (26AWG) after which politeness cutback would be 0).

The case of Figure 2 would be more typical of a residential neighbourhood, where T1 or HDSL services in the same or adjacent binders would be less common. It should be pointed out that if all local subscribers were being served off power-cutback-capable DLCs , the G.lite crosstalk levels would actually be lower due to the 'extra' power cutback and the rates would be improved over those shown in figure 2.

With 10 HDSL disturbers (figure 3), full (1.5Mbps) capacity is possible with 15dB of cutback to loop lengths of 9kft (26AWG) – within CSA loop engineering rules.

With T1 disturbers in adjacent binders (figure 4), the DS power reduction - on a loop of length 12kft - may be restricted to ~6dB to avoid significant capacity hits.

---

[2] The attainable rates assume:

- 4dB SNR margin
- 3dB coding gain
- DS carriers 36-127
- DS receiver noise floor of –136.8dBm/Hz

Where a higher (than 4dB) SNR margin is mandated (by the operator) and/or the DS receiver noise floor exceeds –136.8dBm/Hz (required to meet test cases with DS carriers 36-127) the attainable rate or allowable DS power cutback for a given rate will be reduced.

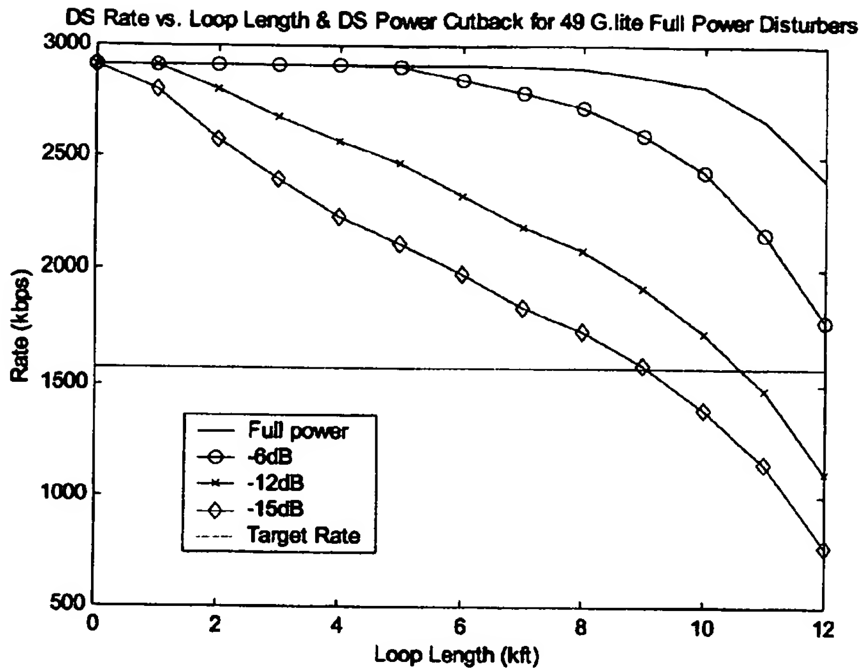

Figure 2 DS Rates with 49 G.lite Disturbers

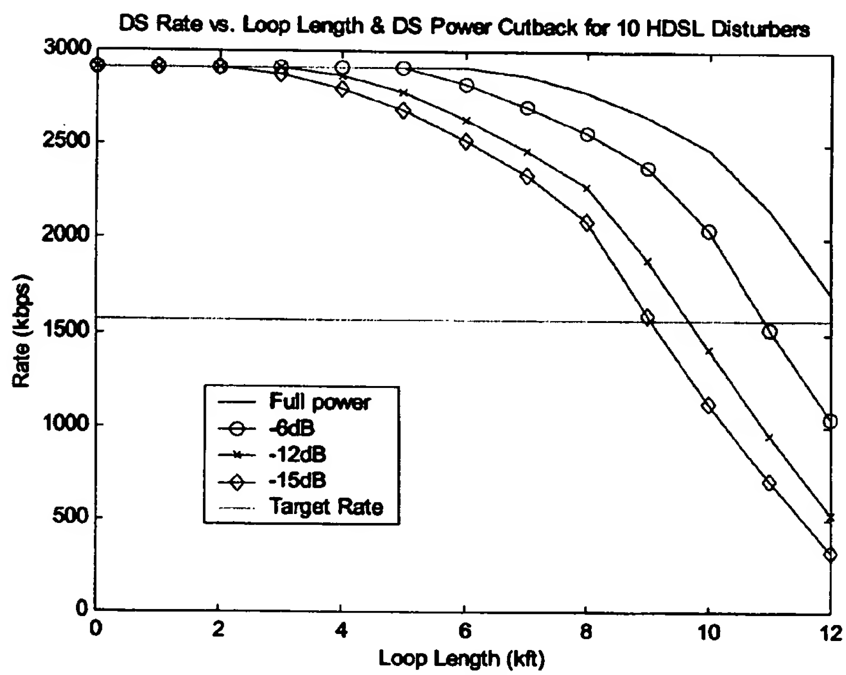

Figure 3 DS Rates with 10 HDSL Disturbers

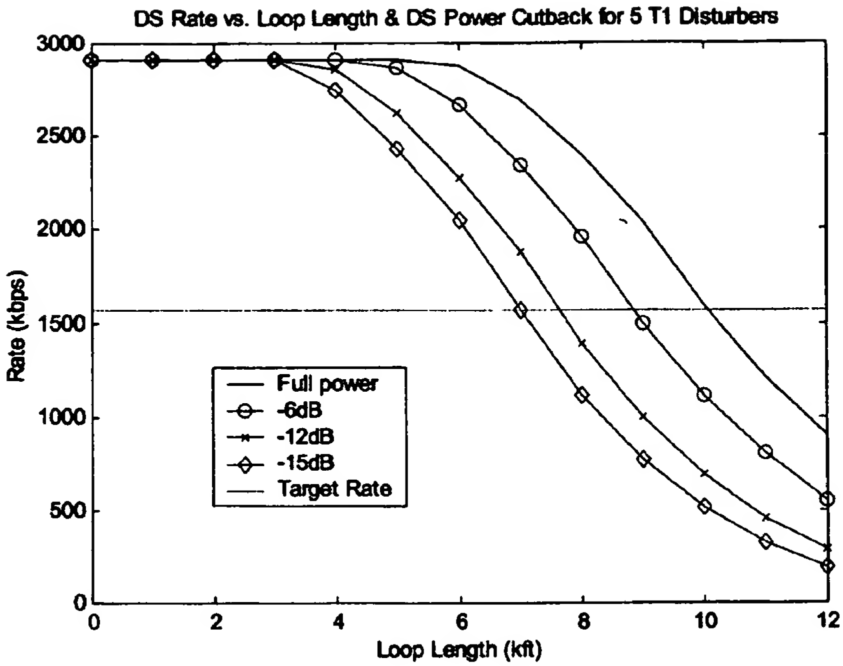

Figure 4 DS Rates with 5 Adj. Binder T1 Disturbers

This invention describes a DS power cutback mechanism that reduces DS transmit power based on a metric of excess SNR margin at the DS receiver. A variety of techniques are described, some of which can be implemented with ATU-Rs compliant to the current draft G.lite standard (ITU-T G.992.2) and some which may be more optimal (faster to initialize and/or maximizing achievable cutback) but require changes to the existing draft standard.

Description of invention

Determination and setting of DS power cutback – where excess SNR margin is available at the DS receiver - is not directly supported in the G.992.x initialization procedure. It can be implemented in a non-standard fashion, for interoperability with G.992.2 ATU-Rs but a more efficient mechanism can be proposed for G.992.2bis.

Non-standardized DS Power Reduction Initialization Procedures

The amount of DS transmit power reduction that can be tolerated for a given capacity target will be a function of the loop and the crosstalk environment which together determine the SNR per receiver sub-carrier. Unfortunately, this information is not

available to the ATU-C. It may be inferred from the Bi's and Gi's sent in R-B&G, near the end of initialization, but it is too late to implement a DS power cutback without impacting the ATU-R, which will need to adapt its receive AGC at a minimum.

In the absence of changes to G.992.2 to address this, there are a variety of approaches to initialization that could permit power savings without significant capacity losses. NOTE: In all cases, the amount of cutback vs. the measured upstream power could be an operator-specified parameter.

A) Reduce DS power based on measured US power alone — PREFERRED

US power is measured at the ATU-C in order to determine the DS power politeness cutback level already.
The politeness cutback is used on very short loops only (~2-3kft 26AWG equivalent), so there is no impact on DS capacity regardless of the crosstalk noise environment at the ATU-R.

It is clearly possible to cutback DS power more aggressively, based on the results shown in figures 1-3,
although the degree to which this is done would depend on the willingness to risk a failure during full-initialization. For example, a cutback of 12dB may be employed where the loop length is estimated to be less than 9kft 26AWG-equivalent. Assuming that the loop length – and equivalent DS loop losses – can be adequately estimated from the US power measured at the ATU-C AND that there are no adjacent binder T1 disturbers, a 12dB cutback should not significantly reduce the DS capacity below ~1.5Mbps.

Unexpected capacity losses may still occur if the ATU-R is not capable of performance meeting the standard test cases because of a noisy front end or if it is not robust against AM ingress noise. In those cases, the cutback could be reduced or eliminated and a second full-initialization triggered.

B) Two passes through full-initialization during initial install or whenever on-hook loop conditions have changed significantly

In this case, the ATU-C forces a second full-initialization once it has received the Bi's, Gi's, performance (SNR) margin and attainable rate from the ATU-R near the end of the first initialization procedure. It can estimate the impact of a cutback from those parameters (if correctly reported) and determine how much cutback can be tolerated. A second initialization procedure would then be repeated with the ATU-C transmitting at its 'new' level.

This process would require twice as long (up to ~20seconds) to initialize the link but would only need to be invoked whenever a change has occurred in the noise environment at the ATU-R (for example, when a
HDSL service is newly deployed).

C) Reduce DS power by excess margin reported by ATU-R

Negotiation of rates on the link occurs over a series of handshakes between the ATU-C and the ATU-R. The ATU-C starts by providing the ATU-R with a list of 4 rate options as well as the required SNR margin to be met. The ATU-R responds to those options indicating the highest rate, if any, that can be supported. It also provides the ATU-C with the average DS loop attenuation and the SNR margin at that rate (should be approximately same at all carriers if the Gi's have been calculated to equalize the SNR margin across all carriers). If the ATU-R is able to support the highest rate, with a SNR margin in excess of the desired margin, then a power reduction of the amount of the difference can be implemented. The difficulty with this approach is that the ATU-R will not necessarily be designed to adapt to a signal level change at this point in the initialization – forcing a new full initialization

D) Over-specify margin in first round of rate negotiations

In the first round of rate negotiations, the ATU-C specifies a minimum SNR margin with an additional 'N' dB of margin. If the ATU-R responds its first reply that it can support a high link rate at that inflated margin – the ATU-C drops its Tx power by 'N' dB and sets the required margin lower by an equivalent amount for the second round of rate negotiations. If the ATU-R cannot support the target link rate with the inflated (by 'N') margin, then the ATU-C does not cutback its transmit power, but still reduces the minimum required SNR margin (by 'N') for the second round of negotiations. This approach suffers the same issue with ATU-R AGC adaptation as in method C above. It also requires that the ATU-R indicate whether it can support any of the options in the first round; the G.992.x standard currently allows the ATU-R to postpone a decision, which would make this approach untenable.

E) Count on ATU-R to signal attainable cutback implicitly in DS Gi's

Finally, it is also possible – but not required – that the ATU-R reduce the Gi's to be used on each of the downstream carriers where there is (otherwise) excess margin . If so, this would – in effect – set the downstream power cutback desired here. As noted, one cannot assume that any vendor's ATU-R will specify the Gi's in that manner (i.e. they may use only to equalize the margin on each carrier, while keeping an overall unnecessarily large SNR margin). In addition, the Gi's are communicated to the ATU-C late in the initialization process. As a result, any change in transmit gain must be exact as there is no

time for the ATU-R to adapt its receiver to a imprecise gain change, before the start of SHOWTIME. In practice,
this forces the signal to be reduced in the digital domain, before the DAC, and so can place excess demands on the DAC dynamic range.

Proposal for DS Power Reduction in G.992.2bis

The amount of DS power reduction that can be tolerated for a given rate and SNR margin requires knowledge of the loop loss crosstalk noise spectral density at each carrier. It is also desirable that the DS power reduction be implemented at the ATU-C in time for the ATU-R to adapt its receiver before SHOWTIME.

The basic objective is to manage the signal power being launched into one end of the loop, so as to provide a signal level, at the receiver end, sufficient to meet the targeted capacity but not higher. In T1.413/G.992.1 or G.992.2, the upstream transmitter cutback may be determined and applied, during link initialization, as follows:

1. The ATU-R measures the DS received[3] noise level (during C-QUIET2)
   - any receiver AGC should be set to maximum gain here to support a noise level measurement of
     sufficient resolution
2. The ATU-R measures the DS received power without cutback (during C-REVERB1) and estimates the average loop attenuation.
3. The ATU-R calculates the amount of DS cutback possible without reducing estimated upstream capacity by more than 2(?)%. The capacity calculations are based on received signal and receiver input referenced noise levels as measured at each upstream carrier frequency and mapped to an equivalent per-carrier *SNRi*. Note that this capacity formulation is based on the measured signal and noise levels, independent of near end echo, timing jitter or mis-equalization effects. An additional 6dB of margin is included in the SNR gap figure used in the capacity calculation to cover these impairments as well as measurement tolerances and to provide a hedge against the introduction of more NEXT sources into the same binder group.

   The line rate capacity estimate, less the FEC overhead, (in bits per symbol period) with no cutback is given by

$$C = \sum_i b_i; \quad b_i = round(\log_2(1 + \frac{SNR_i}{\Gamma})) \quad and\ 2 \le b_i \le b_{\max}$$

   where b_{max} is the maximum # of bits/carrier supported in the ATU-R receiver[4],

[3] A minimum noise floor measurement resolution capability may be specified to ensure that ATU-Rs are capable of supporting a specific DS power cutback for a given DS loop loss and crosstalk noise environment.

[4] Note that the ATU-C transmitter may have a lower value of b_{max}, but that is not known at the ATU-R prior to reception of C-MSG S1, though it could be specified for exchange during G.hs.

and

$$\Gamma = 10^{(9.8+10-3)/10}$$

is the SNR gap for a bit-error-rate of 10^{-7} with 10dB margin and 3dB coding gain, for G.992.2.

The maximum *cutback* (in linear form – corresponding to some multiple of 2dB) is determined such that the capacity with that cutback, C', is greater than or equal to *0.98 x C*, where

$$C' = \sum_i b_i'; \quad b'_i = round(\log_2(1 + \frac{SNR_i / cutback}{\Gamma})) \quad and\ 2 \le b_i \le b_{max}$$

4. The ATU-R encodes the required cutback into a field of R-MSGS1, signaling the correct downstream power cutback to the ATU-C. Allowing for a power cutback range of 0-15dB, in 1dB steps, a 4 bit field is required in R-MSGS-1.

5. After receiving and decoding the message from the ATU-R, the ATU-C applies the power cutback – 500-3500 symbols into C-MEDLEY. The ATU-R is able to adapt its receiver gain, to the new ATU-C transmit level, over this interval prior to completion of C-MEDLEY.

with that excess margin.

5. A method for determining excess SNR margin, for a given target capacity, at the DS receiver early in link initialization and communicating the achievable and required DS power cutback to the DS transmitter in time to implement the cutback and allow the DS receiver to retrain its receive gain before beginning data transmission.

Back-up Material

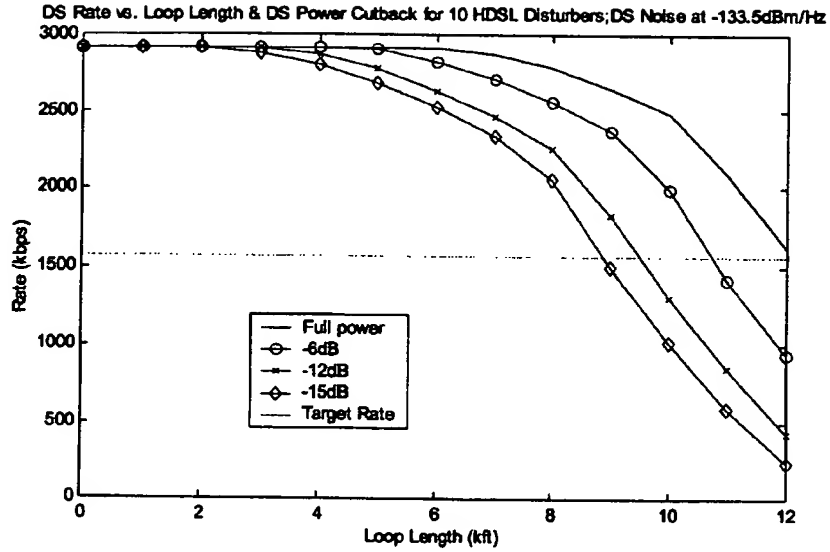

Figure x: With DS noise floor at –133.5dBm/Hz (vs. –136.8)

0

Claims

1. An xDSL xTU-C transciever capable of determining excess SNR margin (at normal transmit levels) on the downstream (DS) link and reducing DS transmit signal power, for a given target DS capacity.

2. A method for estimating excess SNR margin, based on an estimate of loop loss derived from measured received power in the upstream band and a assumed DS receiver input-referenced noise spectral power density.

3. A method for estimating achievable DS power cutback – for a given target DS capacity - based on the reported SNR margin and selected Bi's (bits/carrier 'i') and Gi's (gain per carrier 'i') from the DS receiver – in a multi-carrier system.

4. A method for estimating achievable cutback, based on first over-specifying the required DS SNR margin by an amount equal to the desired cutback, then reducing the DS transmit power for any link that could have supported the target DS capacity